High-Performance Graphics (2023)
J. Bikker and C. Gribble (Guest Editors)

Real-Time Rendering of Glinty Appearances using Distributed Binomial Laws on Anisotropic Grids

T. Deliot and L. Belcour

Intel Corporation



Figure 1: *We render in real-time glittery appearances such as snow, car paints, or christmass ornaments.*

Abstract

In this work, we render in real-time glittery materials caused by discrete flakes on the surface. To achieve this, one has to count the number of flakes reflecting the light towards the camera within every texel covered by a given pixel footprint. To do so, we derive a counting method for arbitrary footprints that, unlike previous work, outputs the correct statistics. We combine this counting method with an anisotropic parameterization of the texture space that reduces the number of texels falling under a pixel footprint. This allows our method to run with both stable performance and $1.5\times$ *to* $5\times$ *faster than the state-of-the-art.*

1. Introduction

While rendering glittery appearances greatly enhances the realism of many human-made surfaces (car paints, molded plastics, ...) and natural ones (stones, water, snow, ...), adding it to a rendering engine comes at very high cost. Indeed, such effects require integrating high-frequency details, either due to counting discrete particles [JHY*14] or integrating high-resolution normal maps [YHJ*14]. In contrast to offline rendering, real-time implementations of glittery surfaces must relax different aspect of what defines such models while at the same time preserving their appearance. This is a rather difficult task, and as such, only a handful of methods managed to add this effect to real-time rendering engines [ZK16, CLS*21].

Specifications for Real-Time. One major constraint is that the distribution of details must remain coherent whatever the distance to the camera is. In a nutshell, if a particular glint exists at a given scale, it must remains when the camera zooms out or rotates. Real-time implementations relax this constraint to only *ensure coherence between two consecutive levels-of-detail* (that is, when the camera footprint quadruples its area).

However, real-time methods still require that the underlying stochastic details *converge towards a smooth microfacet model* when the number of details tends to infinity.

BRDF as a Random Process. Using those specifications, Zirr and Kaplanyan [ZK16] derived the fastest glinty shader so far where the *Normal Distribution Function* (NDF) of a microfacet model [WMLT07] is used as the success probability of multiple coin tosses. By connecting the number of trials to the number of facets within the pixel's footprint, they simulate the randomness of glittery appearances. This process satisfies both constraints as the resulting appearance converges to the continuous microfacet model for an increasing number of trials. They further ensure the coherence between mip-levels by attaching random numbers on the object's surface in texture space.

However, this solution presents two drawbacks: First, it doesn't correctly reproduce the statistics of glittery appearance which *produces artifacts*; Second, the way they attach random numbers to the texture space requires that they manually perform anisotropic filtering. This results in *unstable performance* depending on the viewing conditions and glints density.

Our solution. We use the formulation of Jakob and colleagues [JHY*14] and search to count the number of discrete facets reflecting light towards the camera.

Similar to Zirr and Kaplanyan [ZK16], we randomize the evaluation of the NDF by using a stochastic counting process. However, we link this evaluation to an arrangement of facets at the surface of the object and, hence, *ensure a correct evaluation of the statistics of glints*. More precisely, we measure the proportion of flakes within the pixel footprint as if they were randomly distributed over the surface of the object (see Section 4).

Real-time implementations all rely on discretizing a pixel footprint into overlapping texels at a given level of detail (LOD). When a pixel footprint becomes anisotropic, iterating over all overlapping texels can be a huge computational bottleneck. We show that for glint rendering, we can align the texture parameterizations used for random numbers to pixel footprints and *ensure that, at a given LOD, every footprint maps to a constant number of texels* (see Section 5). This results in both fast and stable performance.

Combined with a *new approximation of the binomial law* and the *use of simplex grids* to fetch random numbers (see Section 6), we obtain a real-time prototype that offers stable performance and is faster that the state-of-the-art methods of Zirr and Kaplanyan [ZK16] and Chermain et al. [CSDD20]. We demonstrate that our prototype is $1.5\times$ to $5\times$ faster than theirs (see Section 7).

2. Previous Work

We split previous work in two categories: offline rendering techniques and real-time methods. We subdivide both categories in two depending on the modeling of the surface geometry: either stochastic particles (glints) or high-frequency normal maps.

2.1. Offline Rendering

Stochastic glints. This first body of work, pioneered by Jakob et al. [JHY*14], represents glints as a discrete set of facets randomly distributed on the surface with orientation following a microfacet model's *Normal Distribution Function* (NDF). The amount of light reflected by the surface depends on the number of reflecting facets in the footprint of the current pixel. This is an expensive 4D hierarchical search. Atanasov et al. [AK16] improve on Jakob et al. [JHY*14] model to build a production ready glint appearance model. Wang et al. [WWH18] accelerate this computation by separating the search of the particle position and orientation. This splits the searching process into two 2D hierarchical searches which is less expensive. They further prefilter the directional statistics to avoid the explicit search in this space.

High-frequency normal maps While this is not restricted to rendering glittery appearances, another body of work models glittery appearances using high-frequency normal maps [YHJ*14]. There, one must integrate over the surface of the pixel footprint, all the normals reflecting light towards the camera. Chermain et al. [CCM20] use a similar paradigm expressed in the microfacet framework, which permits including multiple scattering effects for example [CCM19]. Variants of both methods have been explored to reduce the storage or improve the efficiency [GGN18, CSDD20, WHHY20, CSDD21, AWKK21, DLW*22].

2.2. Real-Time Rendering

Both the stochastic glints and high-frequency normal maps models have been translated into real-time shaders. In our literature review, we will discard methods not reproducing one of those models such as sparkle shaders [BW15, WB16]. While they are extremely efficient, they do not offer realistic behaviour or parameters.

Glint models. Zirr and Kaplanyan [ZK16] follow the discrete microfacets approach. They replace the explicit counting of the number of reflecting facets by a statistical average using a binomial law. Each facet in the pixel footprint is treated as tossed coin with a probability of reflecting light proportional to the target NDF. Since this greatly simplifies the complexity of the computation, they achieve real-time performance. Still, their method presents artifacts and remains too expensive for video games. Alternatives [WDH20] propose different approximations of the counting process but all lead to poorer performance (in the order of $10ms$). Our aim is to provide a model that is *artifact-free*, has *stable* performance and runs *faster* than existing methods.

Patch based models. Inside the formulation of high-frequency normal maps, Chermain et al. [CSDD20] achieve real-time performance but with physically-based, energy-conserving results. They separate the directional 2D statistics as products of 1D statistics. To avoid explicit counting, they precompute a dictionary of filtered 1D NDFs that account for the accumulation of normals within the pixel footprint. This later part introduces several artifacts and limitations. Notably, they cannot use the GGX NDF. We relax the PBR constraint while preserving the appearance and provide a model that is compatible with any NDF.

Scope of our work. In this work, we build on the aggregate counting process of Zirr and Kaplanyan [ZK16]. We show how to maintain the correct counting statistics through levels-of-detail (Section 4), and how to reduce the number of calls to the counting method within a pixel footprint (Section 5). First, we provide an overview of the underlying shading model (Section 3).

3. Real-Time Rendering of Glints

Our shading model starts from the Rendering Equation [Kaj86]:

$$L_{\mathcal{P}} = \int_{\Omega} \rho_{\mathcal{P}}(\omega_i, \omega_o)\, L_i(\omega_o) \cos(\theta_o)\, \mathrm{d}\omega_o, \tag{1}$$

averaging over hemisphere Ω the product of the incoming light L_i with the mean reflectance model over the pixel footprint $\mathcal{P}$:

$$\rho_{\mathcal{P}}(\omega_i, \omega_o) = \frac{F(\langle \omega_i \cdot \mathbf{h} \rangle)\, G(\omega_i, \omega_o)\, D_{\mathcal{P}}(\mathbf{h})}{4 \langle \omega_i \cdot \mathbf{n} \rangle \langle \omega_o \cdot \mathbf{n} \rangle}. \tag{2}$$

Here, $\mathbf{h}$ is the half-vector between the view ω_i and light ω_o vectors, and $F(\cdot)$ and $G(\cdot)$ are respectively the Fresnel and Shadowing/Masking terms. This model behaves like a smooth microfacet model with the NDF, $D_{\mathcal{P}}(\mathbf{h})$, replaced by a counting process $c(\cdot)$:

$$D_{\mathcal{P}}(\mathbf{h}) = \frac{c(N(\mathcal{P}), D(\mathbf{h}))}{N(\mathcal{P})}, \tag{3}$$

with $\mathbb{E}[D_{\mathcal{P}}(\mathbf{h})] = D(\mathbf{h})$, and $N(\mathcal{P})$ is the number of discrete microfacets inside the pixel footprint, and $D(\mathbf{h})$ is a standard NDF function such as GGX or Beckmann. While Jakob et al. [JHY*14] define an explicit counting procedure of the microfacets, it is too computationally intensive for real-time constraints.

Counting Microfacets Instead, we follow the method of Zirr and Kaplanyan [ZK16] and rely on realizations of the binomial law to count the number of reflecting facets:

$$c(N(\mathcal{P}), D(\mathbf{h})) = \frac{D(\mathbf{n})}{R} \times b\left(N(\mathcal{P}), R \times \frac{D(\mathbf{h})}{D(\mathbf{n})}\right). \tag{4}$$

Here, $b(N,p)$ is a sampling of the binomial law for N trials with a probability of success of p, $D(\mathbf{n})$ is the NDF at the shading normal and $R \in [0,1]$ is a user defined ratio. While $D(\mathbf{n})$ ensures that the probability of success remain in $[0,1]$, R offers control over the amount of glints locally. In the following, we use $p = R \times \frac{D(\mathbf{h})}{D(\mathbf{n})}$.

R defines another probability of microfacets being reflecting or non-reflecting for any half-vector. This parameter is *ad-hoc* and we cannot completely link it to a physical quantity. However, in practice, it behaves close to a *microfacet roughness* parameter: low R will reduce the amount of glints and increase their intensity.

Convergence to Smooth Model. In this formulation, $D_{\mathcal{P}}$ converges to the smooth NDF when the number of microfacets in the pixel footprint increases. Indeed, when N increases, the ratio $b(N,p)/N$ converges to the probability $p^{\dagger}$ and all factors cancel out:

$$\lim_{N\to\infty} D_{\mathcal{P}}(\mathbf{h}) = \frac{D(\mathbf{n})}{R} \times \lim_{N\to\infty} \frac{b(N(\mathcal{P}),p)}{N(\mathcal{P})} \tag{5}$$

$$= \frac{D(\mathbf{n})}{R} \times \frac{R \times D(\mathbf{h})}{D(\mathbf{n})} = D(\mathbf{h}). \tag{6}$$

Temporal Coherence To ensure that this counting process is temporally stable, random seeds θ are distributed in a spatial and angular grid and used to evaluate the counting law. Evaluations are bi-linearly interpolated to simulate the spatial and angular life-span of the glint. In their implementation, Zirr and Kaplanyan [ZK16] also use LOD levels to avoid aliasing. This requires interpolating between two spatial grids at a time to achieve a continuous counting law, which they do linearly. Their binomial law is thus:

$$b(N(\mathcal{P}),p) = \text{mix}\left(b_{\overline{\theta}}\left(N \times \overline{\mathcal{P}}, p\right), b_{\underline{\theta}}\left(N \times \underline{\mathcal{P}}, p\right), w\right), \tag{7}$$

where $\text{mix}(a,b,w) = (1-w) \times a + w \times b$, and $\overline{\mathcal{P}}$ and $\underline{\mathcal{P}}$ are the grid cell areas of the top and bottom LOD. This interpolation design has two limitations that prevent it from achieving stable real-time performance with no artifacts:

1. **Incorrect Counting.** Interpolating the outputs of the binomial laws between LOD-levels leads to ghosting artifacts (see Figure 2 (a)) and produces incorrect statistics (see Figure 4 (b)). In Section 4, we detail how interpolating the parameters of the binomials instead of their output solves those issues.
2. **Costly Anisotropic Filtering.** When the pixel footprint covers multiple texels, the counting process becomes expensive. Indeed, the counting law needs to be manually evaluated for each texel under the pixel footprint (Figure 3 (a)). Because of projection, the pixel footprint can become anisotropic with an arbitrary high spread. This leads to uneven performances since any tilted surface w.r.t. the camera will induce some anisotropy. In Section 5, we provide a parameterization for grids of random numbers that adapts to the anisotropy (see Figure 3 (b)).

† since $\mathbb{E}\left[\frac{b(N,p)}{N}\right] = p$, and $\text{Var}\left[\frac{b(N,p)}{N}\right] = \frac{p\times(1-p)}{N}$.

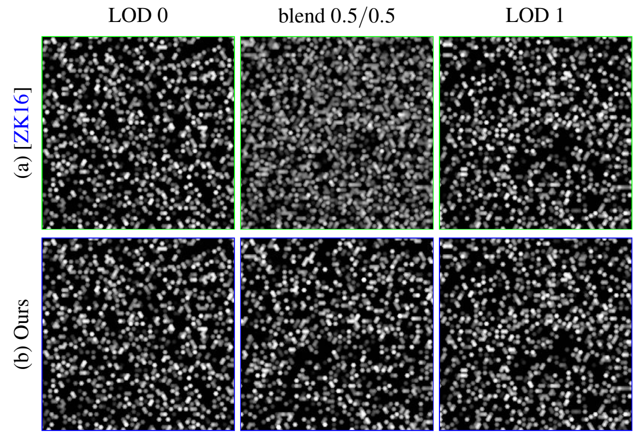


Figure 2: *Interpolating the outcome of the counting law between LOD levels [ZK16] leads to ghosting artifacts (a). Interpolating the parameters of those laws avoids those artifacts (b). In the middle column, we blend with equal weights the right and left insets.*

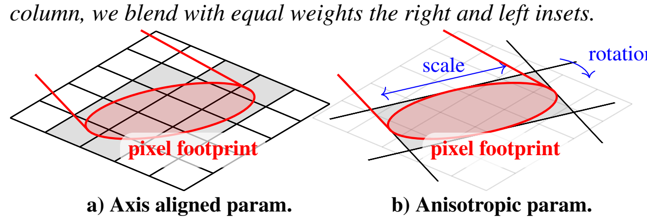


Figure 3: *Projected pixel footprints can be highly anisotropic. When the UV space is axis-aligned, the evaluated number of reflecting glints is distributed over many texels (gray texels - (a)). If we could align the UV-parameterization to the projected pixel footprint, the evaluation would only requires a single texel (b).*

4. Distributed Binomial Laws

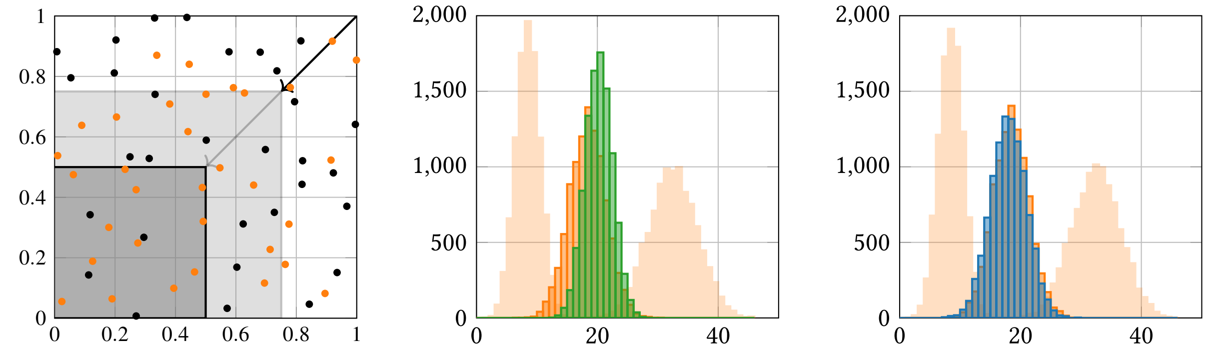
The counting method of Zirr and Kaplanyan [ZK16] discards the spatial location of facets within a footprint. However, each facet has a random position at the surface and is a glint if: it is *within the pixel footprint area*; and, it *reflects light*. Both tests are in fact Bernoulli trials with different probabilities. The combined outcome of those tests is what we call a *spatialized Bernoulli trial*. In this work, we assume that both trials are independent and devise a proper interpolation scheme between LOD levels.

4.1. Spatialized Bernoulli Trials

Let's attach to a random point $\mathbf{x}_i$ on a 2D domain the result of a Bernoulli trial $X_i < p$. We call a spatialized Bernoulli trial the outcome of $\mathbf{x}_i \in \mathcal{A} \cap X_i < p$ for a given area $\mathcal{A}$ and success probability p (see Figure 4 (a)). We further define the *spatialized binomial law* that counts the number of success in N *spatialized* Bernoulli trials. When we use such a counting law to describe a glittery surface, we implicitly take into account the spatial distribution of facets. However, we must to define the spatial distribution of points.

Uniform Distribution. If we randomly distribute the trials at the surface of the object, the probability of having a trial within a given area $\mathcal{A}$ is a Bernoulli trial with $p = \mathcal{A}$. It follows that the number of facets within a pixel footprint of area $\mathcal{P}$ becomes:

$$N(\mathcal{P}) = b(N, \mathcal{P}). \tag{8}$$



a) Bernouilli trials on a plane for $N = 64$ points and $p = \frac{1}{2}$

b) Blending vs Reference for $M = 10000$ trials

c) Distributing vs Reference for $M = 10000$ trials

Figure 4: ***Distributing Binomial Law.*** *We model the counting of glints as a set of Bernoulli trials randomly distributed on a unit plane **(a)**. In this setup, we need to model the distribution of successes (orange dots) in a given window (gray rectangles). Previous works such as Zirr and Kaplanyan [ZK16] blend between two binomials representing successive mip-levels (**b**) light orange). They obtain a distribution (in green) that does not account for the spatial distribution. We introduce the **Distributed Binomial law (c)** that reproduces the correct statistics.*

Stratified Distribution. When the trials are stratified at the surface of the object (that is, randomly jittered inside a regular grid), the average number of trials for a given area $\mathcal{A}$ is exactly $N_{\mathcal{A}} = N \times \mathcal{A}$. Thus, the number of facets within an area $\mathcal{P}$ becomes:

$$N(\mathcal{P}) = \lfloor N \times \mathcal{P} \rceil. \tag{9}$$

Uniform or Stratified? We described two ways of distributing the Bernoulli trials at the surface of the object. Both relate to ways of modeling glints: a discrete set of random facets [JHY*14], or a square texture with exactly one random facet per texel (akin to normal maps [YHJ*14]). We opted for the later as the associated counting law was less sensitive to numerical precision.

4.2. Blending of Distributed Binomial Laws

To express the counting law between LOD levels, we aim at reproducing the statistics of the stratified distribution of Bernoulli trials for intermediate footprints between the LOD levels. For a given footprint area $\mathcal{P}$ between a top LOD level of area $\overline{\mathcal{P}}$ and seed $\overline{\theta}$, and bottom LOD level of area $\underline{\mathcal{P}}$ and seed $\underline{\theta}$, we evaluate the number of glints as:

$$b(N \times \mathcal{P}, p) = b_{\overline{\theta}}\left(\overline{w} \times N \times \overline{\mathcal{P}}, p\right) + b_{\underline{\theta}}\left(\underline{w} \times N \times \underline{\mathcal{P}}, p\right), \tag{10}$$

where $\overline{w} = {}^{\mathcal{P} - \underline{\mathcal{P}}}/_{\overline{\mathcal{P}} - \underline{\mathcal{P}}}$ and $\underline{w} = {}^{\overline{\mathcal{P}} - \mathcal{P}}/_{\overline{\mathcal{P}} - \underline{\mathcal{P}}}$. Intuitively, this blending subdivides the area $\mathcal{P}$ in two: $\mathcal{P} = \overline{w} \times \overline{\mathcal{P}} + \underline{w} \times \underline{\mathcal{P}}$, with $\underline{w} = 1, \overline{w} = 0$ when $\mathcal{P} = \underline{\mathcal{P}}$, and $\overline{w} = 1, \underline{w} = 0$ when $\mathcal{P} = \overline{\mathcal{P}}$. Then, the binomial for the whole area $\mathcal{P}$ is *distributed* as the sum of binomials in each sub-area. This is possible thanks to the additivity of the binomial law: $b(N_1, p) + b(N_2, p) = b(N_1 + N_2, p)$. Different distributions are possible (for example distributing over more than 2 areas) but follow the same procedure.

This scheme provides the correct statistics for stratified grids of Bernoulli trials (see Figure 4) and avoids ghosting artifacts (see Figure 2). While this corrects the glints' statistics, it still requires a manual anisotropic filtering. In the following Section, we devise a method to always bring the evaluation back to the isotropic case.

5. Anisotropic Virtual Grids

In our method, similar to previous works [ZK16, CSDD20], multiple virtual grids of random numbers exist at the surface of the object. For each pixel, the grids with texels best matching the pixel footprint are used to evaluate the counting law. In Zirr and Kaplanyan [ZK16] the grids follow a traditional mip-mapping parameterization: uniform grids of square texels (see Figure 5 (a)). This can be done by uniformly scaling UV coordinates and keeping equally sized grids:

$$\theta(\mathbf{uv}') \quad \text{with } \mathbf{uv}' = \mathcal{S}_s(\mathbf{uv}), \tag{11}$$

where $\mathcal{S}_s$ is the isotropic scale. Since texels are square, distributing the binomial law on such grid approximates the counting of an isotropic filter. Counting anisotropic filter requires accumulating the results of counting in all the texels covered by the filter. We extend the idea of changing the UV parameterization to include non-axis-aligned and non-isotropic grids: by covering the object with grids that are linear transformations of the isotropic grid, we can emulate anisotropic filters (see Figure 5 (b-c)).

5.1. Anisotropic Parameterizations

To ensure that each pixel filter corresponds to one texel, we can transform the UV coordinates using a rotation and a ratio to align the texels with the projected pixel footprint:

$$\mathbf{uv}' = \mathcal{R}_{\phi}\left(\mathcal{S}_{1,r}\left(\mathcal{S}_s\left(\mathbf{uv}\right)\right)\right) \tag{12}$$

where $\mathcal{S}_{1,r}$, and $\mathcal{R}_{\phi}$ are respectively the isotropic scale s and anisotropic scale r for the v coordinates, and the rotation of angle ϕ. In practice, we need to use a discrete set of scales and rotations, which produce a countable set of grids. To find the closest fitting grids for a given pixel filter, we must evaluate its rotation and relative scales.

We start by projecting the pixel footprint on the surface of the object and extract its associated ellipse's major and minor axis $\mathbf{v}^0, \mathbf{v}^1$ in UV space. From this vector, we obtain a rotation $\phi = \text{atan}(\mathbf{v}^0_x, \mathbf{v}^0_y)$ and a ratio $r = {}^{||\mathbf{v}^0||}/_{||\mathbf{v}^1||}$. In our implementation, we use screen-space derivatives to obtain the minor/major axis.

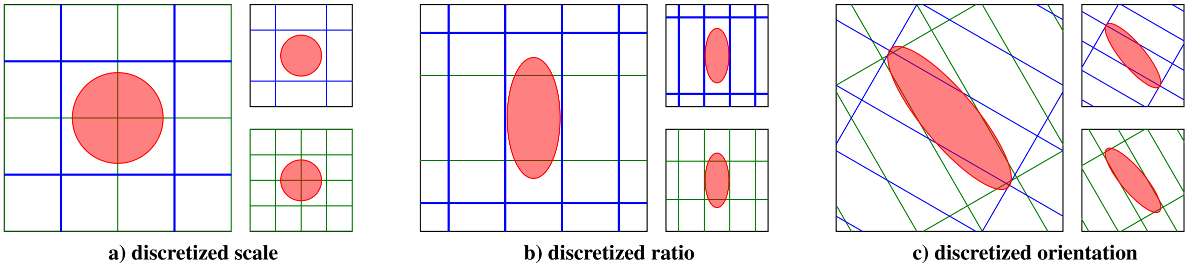


Figure 5: *We attach random numbers on a discrete set of grids. For a given pixel filter (in red), we must select the closest grids that encompass the pixel filter's area (grids in blue and green). Because previous works only account for uniform grids with power of two ratios (a), the evaluation of a grid cell can only account for an isotropic footprint. We add the support of anisotropic footprints using grids with a discrete set of ratios (b) and a discrete set of orientations (c). Note that we aligned the pixel filter with the grids for visualization purposes.*

5.2. Interpolating Different Parameterizations

In the case of the isotropic parameterization, every pixel footprint's area is enclosed by the cell areas of the isotropic grids:

$$\underline{\mathcal{P}} = \text{pow}\left(2, \lfloor \log_2(\mathcal{P}) \rfloor\right), \text{ and } \overline{\mathcal{P}} = \text{pow}\left(2, \lceil \log_2(\mathcal{P}) \rceil\right) \quad (13)$$

Similarly, the ratio and orientation of a pixel footprint are enclosed by discrete ratios and angles. This means that $2^3 = 8$ different texels enclose the pixel footprint. We use the blending formula for the distributed binomial law, Equation 10, over those 8 areas. To do so, we assign to each cell a number of facets N_i proportional to its area. We use trilinear weights w_i to distribute the pixel footprint area over each cell's area. We obtain the final distributed binomial value by summing over all cells: $b(N, p) = \sum_i b_{\theta_i}(w_i \times N_i, p)$.

We adapt the discretization of the orientation ϕ with respect to the anisotropy ratio r. This is to account for the fact that, for isotropic parameterizations ($r = 1$) the orientation of the filter should not alter the random numbers θ. We show in Figure 6 the discretization of the orientation with respect to the anisotropy ratio.

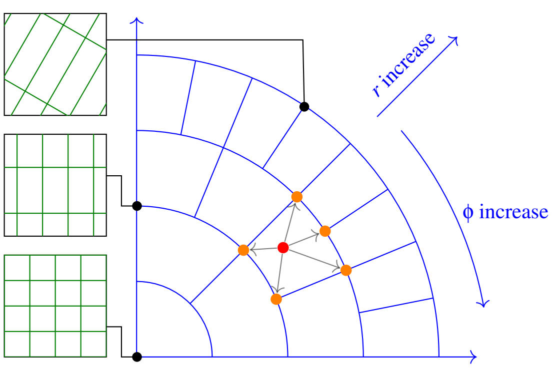


Figure 6: *We discretize the change in orientation with respect to the anisotropy ratio: when no anisotropy is required, a single isotopic grid accounts for orientation of the pixel filter. We add one orientation per level of anisotropy ratio. Thus, a pixel filter (red) distributes its binomial law on 5 grids (orange).*

Angular and Spatial Interpolation. Virtual grids account for levels-of-detail filtering, to prevent aliasing. In-between these, it is important to *distribute* the counting law. However, to reproduce the angular lifespan of individual glints (to account for the smooth transition due to the micro-roughness) or their footprint, we must *interpolate* between *realizations* of the counting law on the grid. For the angular behaviour, we bi-linearly blend evaluations of the same binomial law:

$$b_{\theta_i}(w_i \times N_i, p) = \sum v_j \, b_{\theta_{i,j}}(w_i \times N_i, p). \quad (14)$$

where v_j are the bi-linear weights for the four closest texel of an hemispherical grid where $\theta_{i,j}$ are the associated random seeds. We use the orthographic projection of the half-vector onto the plane to index this grid, and subdivide it to account for the glints' micro-roughness.

Thanks to all this, we obtain an *artifact-free* and *stable* shader model for glints. We found that we could further improve its performance during implementation and detail how in the next Section.

6. Implementation Details

Our method heavily relies on evaluations of the binomial law. Indeed, a naive implementation of our method would lead to $10 \times 4 \times 4 = 160$ evaluations of the binomial per pixel (2×5 anisotropic grids with linear blending of 4 spatial texels and 4 angular texels each). We reduce the number of calls to the binomial law by using tetrahedral discretizations (Subsection 6.1) and a simplex spatial grid (Subsection 6.2), as well as the cost of evaluating the binomial law by using a gated approximation (Subsection 6.3)

6.1. Optimized Anisotropic Parameterization Blending

As presented in Section 5, our anisotropic grids interpolation scheme uses a 3D grid. Since we increase the discretization of orientation at each level of the anisotropy ratio, each cell of the parameterization comprise 3 different orientations and requires 10 vertices for distribution of the binomial law (5 for ratio and orientation, times 2 for scale). We reduce this cost by tessellating this grid using tetrahedra (see Figure 7).

In the $r - \phi$ dimension, each cell is a pentagon. We split these into three triangles (see Figure 7 (left)). Along the s dimension, each triangle becomes a prism that we can further split into three tetrahedra. The overall grid cell, can thus be decomposed into nine different tetrahedra. At run-time, we first identify which tetrahedron encompasses the pixel filter and use its barycentric coordinates to distribute the binomial law. With this, we reduce the number of binomial evaluation from 160 down to $4 \times 4 \times 4 = 64$ per pixel.

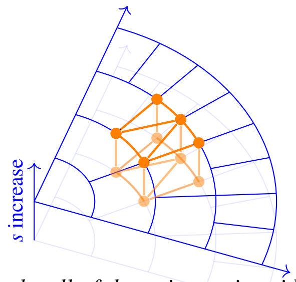


Figure 7: *Each cell of the anisotropic grid can be split into nine different tetrahedra. Using this tessellation, we reduce the number of binomial evaluation and improve our shader's performance.*

6.2. Simplex Surface Space Grids

We can further decrease the amount of binomial samples by applying the same logic to the anisotropic grids themselves: instead of looking at grids with 4 vertices per cell, we tessellate these grids into simplex grids (grids of triangles). While sampling the surface parameterization, we discretize the surface space into a 2D simplex grid instead of a regular square grid (Figure 8) following the simplex noise method [Per01].

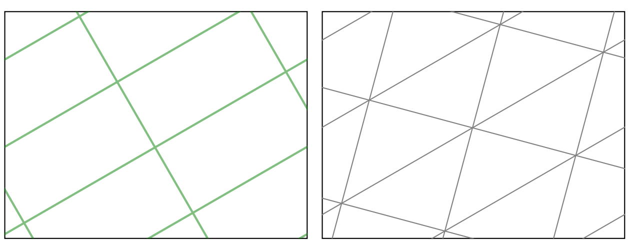


Figure 8: ***Surface Simplex.*** *We parameterize the surface of anisotropic grids (in green) using a simplex grid (in gray). It allows reducing the number of spatial samples from 4 to 3 per pixel.*

Evaluation of a pixel in a triangular grid requires 3 evaluations with barycentric interpolation, instead of 4 bi-linearly interpolated samples. This also modifies the appearance of the glints: square grids with bi-linear interpolation produce square glints, but triangular grids with barycentric interpolation produce hexagonal glints. This results in a rounder appearance which we believe is more natural.

With this, we bring the total down to $4 \times 3 \times 4 = 48$ binomial evaluations per pixel. While we could also reduce the complexity by using a simplex grid in the angular domain, we found that the overall gain was minimal due to a bigger complexity of handling a simplex grid in this space.

6.3. Approximation of the Binomial Law

With this many evaluations still required per pixel, the performance of sampling the binomial law becomes critical. We found that the approximation of Zirr and Kaplanyan [ZK16] was too costly for our use case.

Classical Binomial Approximation. Indeed, Zirr and Kaplanyan use a Gaussian approximation of the binomial law using its first two moments (mean and variance) for which analytical forms are known: $\mathbb{E}[b(N,p)] = N \times p$ and $\mathrm{Var}[b(N,p)] = N \times p \times (1-p)$.

However, to better approximate the discrete nature of the binomial at low number of trials N or probability p, they opt to manually evaluate the sum of the Bernoulli trials:

$$b(N,p) \simeq \begin{cases} \lfloor \mathcal{N}\left(\theta_0; \mu, \sigma^2\right) \rfloor & \text{if } N \times p > 5, \\ \sum_{k=1}^{N} H(\theta_k - p) & \text{else.} \end{cases} \quad (15)$$

Here, $H(x)$ is the Heaviside function. This form impacts the stability of the performance due to the varying sum length. To remedy this, we introduce a fast binomial approximation that always uses the same form, independent of the parameters N and p.

Gated Gaussian Binomial Approximation. We devise a different approximation to the binomial distribution which only requires two random numbers to sample and no looping trials. Our approximation use the Gaussian approximation of the binomial law but aims at improving its statistics for low N and p parameters through the use of a single additional *gating* Bernoulli trial.

We first compute the probability of having at least one success in a binomial distribution $b(N,p)$:

$$P_{\geq 1} = 1 - (1-p)^N, \quad (16)$$

and compute the parameters μ, σ of a Gaussian approximation $\mathcal{N}$ of the binomial distribution for one less sample, $b(N-1,p)$:

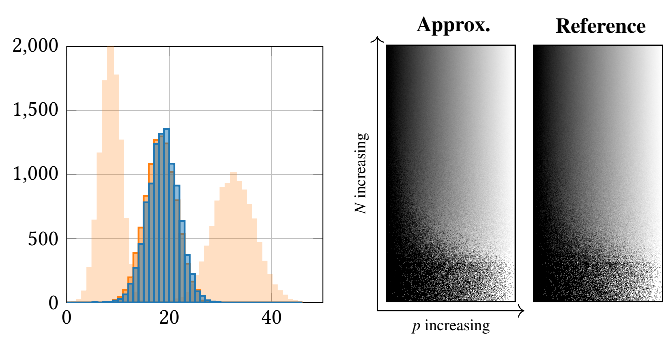
$$\mu = (N-1) \times p, \quad \sigma = \sqrt{(N-1) \times p \times (1-p)}. \quad (17)$$

Finally, we fetch two uniform random numbers θ_0, θ_1 and discretize the continuous Gaussian distribution $\mathcal{N}$ by *gating* it using a single Bernoulli trial with $p = P_{\geq 1}$:

$$b(N,p) \approx \lfloor H(\theta_0 - P_{\geq 1}) \times (1 + \mathcal{N}(\theta_1, \mu, \sigma)) \rfloor. \quad (18)$$

We validated that this approximation follows the reference binomial distribution (see Figure 9), even while blending using our distributed binomial law method, as described in Section 4.

As this approximation only requires two random numbers to sample, we can further speed up run-time performance by fetching them from a small texture of random numbers that we precompute once at startup, in a few milliseconds on the GPU.



a) Histogram with Approx. same configuration as Figure 4 **b) Realizations of Binomial Laws**

Figure 9: ***Approximation of the Binomial Law.*** *A gated Gaussian approximation is visually close to the binomial distribution (b) and works well with the distributed binomial law (a).*

7. Results

We implemented our method in the Unity3D game engine, along with the methods of Zirr and Kaplanyan [ZK16] and of Chermain et al. [CSDD20] using their provided source code. We implemented our method in both Unity's *High Definition Render Pipeline* to render Figure 1, and Unity's built-in Forward renderer for every other Figure and for performance measurements. There, we removed any texture and disabled post-processing (except for a *Bloom* filter to visualize high glint intensities). Both implementations are provided in our supplemental material and showcased in our video.

7.1. Performance

We benchmarked all three methods in different rendering scenarios: the PLANE scene where a plane is uniformly filling the screen with glinty material; the HEIGHTMAP scene where the glinty plane is displaced with a high-frequency heightmap; the SPONZA scene with all objects glinty (shown in Figure 11); the SPHERE and MONKEY scenes with objects in the middle of the screen, as shown in Figures 12 and 13. To measure the impact of anisotropy, we report different viewing configurations (facing at 90° and grazing at 25°) for the PLANE and Heightmap scenes. For all scenes, we use the same Beckmann NDF with a roughness of $\alpha = 1.0$. To test the impact of the binomial/filtering evaluation on performances, we measured timings while varying the *microfacet density* for all methods and report in Table 1 the minimum and maximum timings that we measured, on an NVIDIA RTX 2080 GPU with a 1920×1080 image resolution.

Scene	No Glints	Ours	[ZK16]	[CSDD20]
SPHERE	0.49	**0.79 - 0.82**	1.15 - 1.41	1.19 - 1.27
MONKEY	0.49	**0.67 - 0.68**	0.93 - 1.09	1.01 - 1.04
PLANE 90°	0.49	**1.15 - 1.25**	1.40 - 1.83	1.36 - 1.70
PLANE 25°	0.49	**1.15 - 1.26**	2.50 - 3.18	3.25 - 4.61
HEIGHTMAP 90°	0.51	**1.52 - 1.63**	1.93 - 2.48	2.19 - 2.28
HEIGHTMAP 25°	0.51	**1.44 - 1.54**	3.29 - 4.15	4.27 - 4.83
SPONZA	0.53	**1.32 - 1.38**	3.84 - 5.03	5.06 - 5.90

Table 1: *Rendering times in milliseconds per frame in our minimalist test setup. We compare the performance of a baseline Beckmann NDF function, our glinty NDF, and state-of-the-art methods [ZK16, CSDD20]. For each test, we report the min and max timings observed while varying the microfacet density parameter.*

Analysis. Two different aspects of our method allow it to achieve much more stable performance than other methods across the range of rendering scenarios: *our anisotropic grids allow constant time evaluation* for all possible pixel footprints. Both Zirr and Kaplanyan [ZK16] and Chermain et al. [CSDD20] rely on manually sampling the anisotropic footprint, with a varying impact on performance for each pixel.

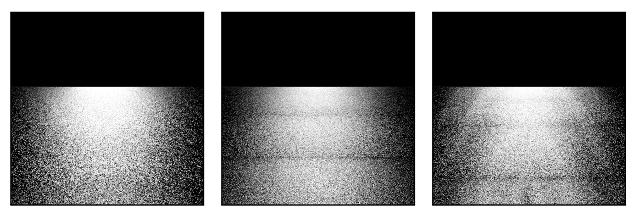
Additionally, *our binomial approximation allows constant time evaluation* for all possible binomial parameters N and p. Zirr and Kaplanyan [ZK16] rely on manually looping Bernoulli trials (up to $N = 16$ in their implementation) resulting in a variable evaluation cost, peaking right before switching to a cheaper Gaussian approximation.

The performance hit related to both aspects are further emphasized by the SIMD architecture of GPUs. With threads running in groups (*waves* or *warps*), a single pixel with large anisotropy (or large amount of Bernoulli trials) stalls the entire thread group, spreading the worst case scenario performance across many pixels. The impact of this can be seen in the performance hit between our PLANE and HEIGHTMAP scenarios at a grazing angle, with the former concentrating the large anisotropy zones together at the top of the screen and the latter scattering pixels of large anisotropy across the entire screen.

The stable performance provided by our method is a highly desirable property for use in real-time production environments. In practice, the performance to consider for a real-time rendering method is its worst-case performance, as interactive applications cannot afford potential slowdowns in their tight frametime constraints. Furthermore, in our testing, we found that the more *real use case* scenario of the SPONZA scene, displaying a wide variety of viewing configurations across all pixels, provided the worst-case scenario and thus showed the most benefit from using our method.

7.2. Renderings

Blending Artifacts. Previous methods [ZK16, CSDD20] display ghosting artifacts (see Figure 10). Those are due to incorrectly blending between either counting laws or discrete NDFs. Our glinty NDF generated using the distributed binomial laws introduced in Section 4 doesn't suffer from this issue.



a) Ours b) [ZK16] c) [CSDD20]

Figure 10: ***Blending artifacts.*** *Our method (a) does not suffer from blending artifacts that one can see on previous works (b-c). They experience such artifact due to linear blending between LOD levels. Such blending does not preserve the counting of glints.*

Anisotropic Footprints. In Figure 11, we show a view point of the SPONZA scene that results in pixel footprints with anisotropy ratios higher than $\times 16$ (white area). With the maximum allowed anisotropy set to $\times 16$ in both [ZK16] and [CSDD20], their evaluation cannot account for higher ratio and, hence, it implicitly increases the size of the footprint isotropically. This results in glints becoming larger on the screen in an unrealistic manner. Increasing the maximum allowed anisotropy would come at an exponentially increasing performance cost and timings variability. Using the anisotropic grid parameterizations introduced in Section 5 avoids these issues by directly sampling the surface using an appropriate parameterization.

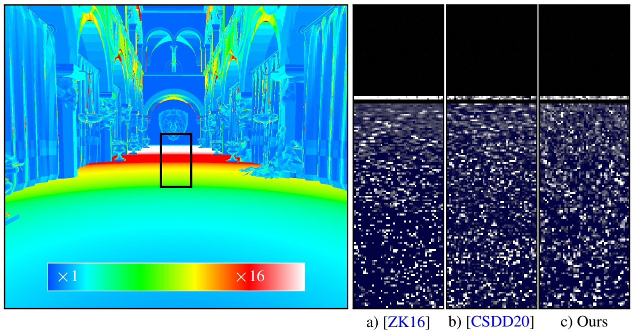


Figure 11: ***Artifacts from anisotropy clamping.*** *In this view of the* SPONZA *scene (left), we measured that the max anisotropy used in previous work is reached on the ground floor. This clamping of the max anisotropy leads to glints stretching in an unrealistic manner (a-b) that is not present in our work (c).*

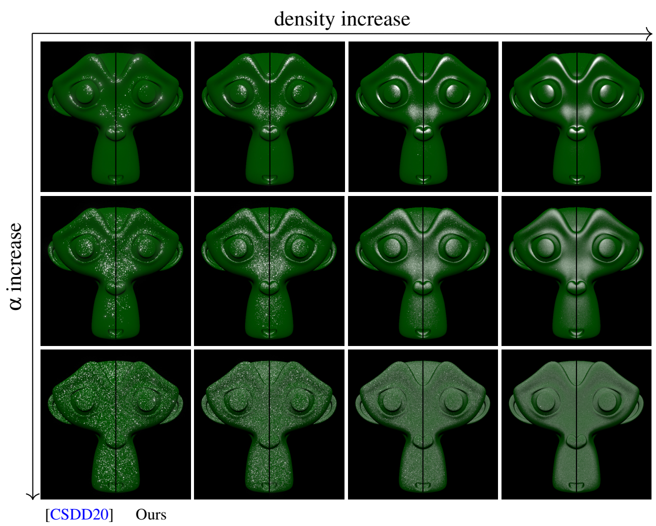


Figure 12: ***Visual comparison with Chermain et al. [CSDD20].*** *We compare our model (right of each image) and the one of Chermain et al. [CSDD20] (left of each image) for different roughnesses and microfacet densities. We validate that our model, while not fully physically based, is visually close to one that is.*

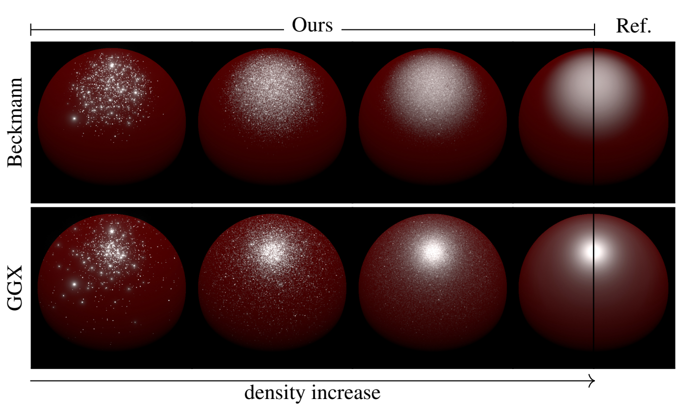


Figure 13: ***Convergence to target NDF.*** *When the density of facets increases, our model converges to any target NDF (last column). Here, we show both Beckmann (top row) and GGX (bottom row).*

Glinty Appearance. In Figure 12, we compare our method with the results of Chermain et al. [CSDD20]. As they produce a *physically-based* BRDF model that is both energy conserving, multi-scale, and continuous, we treat it as the visual reference for our method here.

We use the default parameters for the precomputed dictionary of NDFs in [CSDD20] and use $R = 0.034$ for our method as it closely matches it. We then vary the target Beckmann macro-scale roughness and microfacet density parameters, and capture the result of both methods with matching densities.

Our results show that we can perfectly emulate the physically-based glinty NDF in fixed images using our method. But glints feature very important temporal properties, such as smooth transitions when the viewing and/or lighting angles change. We refer to our supplemental video for temporal comparisons.

Convergence to the target NDF. Both [ZK16] and [CSDD20] limit themselves to simulating glinty NDFs converging to the Beckmann NDF, the former due to requiring analytical convolution of the NDF, the latter due to needing a separable NDF to reconstruct 2D NDFs at run-time from stored 1D NDFs.

We do not share those limitations: the formulation of our NDF as a stochastic counting process shown in Section 3 trivially converges to the target NDF $D(\mathbf{h})$. Intuitively, our formulation is solely redistributing an input energy stochastically, regardless of the function used to compute it. Our method can thus be used to achieve glinty appearances converging to any desired NDF function.

We illustrate this by targeting both the Beckmann and GGX distributions, in Figure 13, across varying microfacet densities. The GGX distribution has become ubiquitous in the real-time rendering industry and is the standard used in modern rendering pipeline materials, such as Unity's *High Definition Render Pipeline* that we use to render the scene shown in Figure 1. Hence, compatibility with GGX will help the wider adoption of real-time glinty materials.

8. Limitations & Discussion

While our model can be a direct replacement of previous works for the rendering of glints in real time, there are some cases where alternative methods would be preferable or more research is required.

Anisotropic glints. While nothing prevents us from using an anisotropic NDF $D(\mathbf{h})$, we did not tackle the problem of rendering *anisotropic glints*. Previous works link the NDF anisotropy to the micro-roughness's anisotropy as it simulates brushed metal. They do so by stretching the UV coordinates (similar to our anisotropic grids). Adding such effect in our work would create with a conflict between the footprint's anisotropy and the glint's anisotropy. For such a case, we argue that a scratch model would be more appropriate [RGB16].

Grids rely on UV Mapping. Similar to other methods, our model requires a UV layout for both the definition of the pixel footprint (via screen-space derivatives) and the evaluation of the anisotropic grid. Unfortunately, the visual quality of the glittery appearance depends on the quality of the UV mapping. A mapping with many discontinuities will result in visible artifacts.

Non-physically-based. Our model to render glint is not physically-based and would, for example, not pass the white furnace test. We stochastically distribute the evaluation of the NDF with potentially non-energy-conserving pixels. As shown in our results, it is not a problem in practice: we obtain renders similar to the PBR model of Chermain et al. [CSDD20] in a procedural way. On the one hand, this frees us from some limitations introduced by their precomputed NDF dictionary (no fewer than four facets at once, micro-roughness constrained by macro-scale roughness). On the other hand, since we do not generate an explicit glinty NDF, we cannot importance sample it. Hence, our solution is not applicable for path tracing for example.

Area lighting and Image Based Lighting. In Section 3 we have defined the counting process w.r.t. the NDF. Unfortunately, this does not account for the light source. As a result, we can only render punctual light with our method. Handling extended light sources would require including the integration with the light source inside the counting process.

Cost per light. Since the counting process depends on the half-vector, our shader needs to be evaluated separately for each light. We cannot share the evaluations of the binomial law between lights. Hence, our shader cost is linear w.r.t. the number of lights.

Both restrictions (extended lights and per light evaluation) are, to us, the last major hurdles preventing the use of our method in production. While some approximations can be done, for example distributing the result of the light integral instead of the NDF, we leave them for future work.

9. Conclusion

We have presented a new method to render glittery appearances in real-time. Thanks to a correct handling of the statistical counting of glints in pixel footprints and a new grid mechanism emulating anisotropic filtering, we achieve *stable* performances and improve by up to 5× the performance of previous work while reducing the number of visual artifacts. While our shader is still costly compared to smooth PBR ones, we believe it will help the adoption of glittery appearances in real-time graphics such as video-games.

References

[AK16] ATANASOV A., KOYLAZOV V.: A practical stochastic algorithm for rendering mirror-like flakes. In *ACM SIGGRAPH 2016 Talks*. 2016, pp. 1–2. 2

[AWKK21] ATANASOV A., WILKIE A., KOYLAZOV V., KŘIVÁNEK J.: A multiscale microfacet model based on inverse bin mapping. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 103–113. 2

[BW15] BOWLES H., WANG B.: Sparkly but not too sparkly! anti-aliasing a procedural sparkle effect. In *Siggraph 2015 Advances in Real-time Rendering in Games* (2015). 2

[CCM19] CHERMAIN X., CLAUX F., MÉRILLOU S.: Glint rendering based on a multiple-scattering patch brdf. In *Computer Graphics Forum* (2019), vol. 38, Wiley Online Library, pp. 27–37. 2

[CCM20] CHERMAIN X., CLAUX F., MÉRILLOU S.: A microfacet-based brdf for the accurate and efficient rendering of high-definition specular normal maps. *The Visual Computer 36* (2020), 267–277. 2

[CLS*21] CHERMAIN X., LUCAS S., SAUVAGE B., DISCHLER J.-M., DACHSBACHER C.: Real-time geometric glint anti-aliasing with normal map filtering. *Proceedings of the ACM on Computer Graphics and Interactive Techniques 4*, 1 (2021), 1–16. 1

[CSDD20] CHERMAIN X., SAUVAGE B., DISCHLER J.-M., DACHSBACHER C.: Procedural physically based brdf for real-time rendering of glints. In *Computer Graphics Forum* (2020), vol. 39, Wiley Online Library, pp. 243–253. 2, 4, 7, 8, 9

[CSDD21] CHERMAIN X., SAUVAGE B., DISCHLER J.-M., DACHSBACHER C.: Importance sampling of glittering bsdfs based on finite mixture distributions. In *EGSR (DL)* (2021), pp. 45–53. 2

[DLW*22] DENG H., LIU Y., WANG B., YANG J., MA L., HOLZSCHUCH N., YAN L.-Q.: Constant-cost spatio-angular prefiltering of glinty appearance using tensor decomposition. *ACM Transactions on Graphics (TOG) 41*, 2 (2022), 1–17. 2

[GGN18] GAMBOA L. E., GUERTIN J.-P., NOWROUZEZAHRAI D.: Scalable appearance filtering for complex lighting effects. *ACM Trans. Graph. 37*, 6 (2018), 277–1. 2

[JHY*14] JAKOB W., HAŠAN M., YAN L.-Q., LAWRENCE J., RAMAMOORTHI R., MARSCHNER S.: Discrete stochastic microfacet models. *ACM Transactions on Graphics (TOG) 33*, 4 (2014), 1–10. 1, 2, 4

[Kaj86] KAJIYA J. T.: The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques* (1986), pp. 143–150. 2

[Per01] PERLIN K.: Noise hardware. In *Real-Time Shading SIGGRAPH Course Notes* (2001). 6

[RGB16] RAYMOND B., GUENNEBAUD G., BARLA P.: Multi-scale rendering of scratched materials using a structured sv-brdf model. *ACM Transactions on Graphics (TOG) 35*, 4 (2016), 1–11. 8

[WB16] WANG B., BOWLES H.: A robust and flexible real-time sparkle effect. In *EGSR 2016 E&I-Eurographics Symposium on Rendering-Experimental Ideas & Implementations* (2016), The Eurographics Association, pp. 49–54. 2

[WDH20] WANG B., DENG H., HOLZSCHUCH N.: Real-time glints rendering with pre-filtered discrete stochastic microfacets. In *Computer Graphics Forum* (2020), vol. 39, Wiley Online Library, pp. 144–154. 2

[WHHY20] WANG B., HAŠAN M., HOLZSCHUCH N., YAN L.-Q.: Example-based microstructure rendering with constant storage. *ACM Transactions on Graphics (TOG) 39*, 5 (2020), 1–12. 2

[WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques* (2007), pp. 195–206. 1

[WWH18] WANG B., WANG L., HOLZSCHUCH N.: Fast global illumination with discrete stochastic microfacets using a filterable model. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 55–64. 2

[YHJ*14] YAN L.-Q., HAŠAN M., JAKOB W., LAWRENCE J., MARSCHNER S., RAMAMOORTHI R.: Rendering glints on high-resolution normal-mapped specular surfaces. *ACM Transactions on Graphics (TOG) 33*, 4 (2014), 1–9. 1, 2, 4

[ZK16] ZIRR T., KAPLANYAN A. S.: Real-time rendering of procedural multiscale materials. In *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (2016), pp. 139–148. 1, 2, 3, 4, 6, 7, 8